

Information Security as Strategic (In)effectivity

Wojciech Jamroga¹ and Masoud Tabatabaei²

¹ Institute of Computer Science, Polish Academy of Sciences

² Interdisciplinary Centre for Security and Trust, University of Luxembourg
w.jamroga@ipipan.waw.pl, masoud.tabatabaei@uni.lu

Abstract. Security of information flow is commonly understood as preventing *any* information leakage, regardless of how grave or harmless consequences the leakage can have. Even in models where each piece of information is classified as either sensitive or insensitive, the classification is “hardwired” and given as a parameter of the analysis, rather than derived from more fundamental features of the system. In this work, we suggest that information security is not a goal in itself, but rather a means of preventing potential attackers from compromising the correct behavior of the system. To formalize this, we first show how two information flows can be compared by looking at the adversary’s ability to harm the system. Then, we propose that the information flow in a system is *effectively information-secure* if it does not allow for more harm than its idealized variant based on the classical notion of noninterference.

1 Introduction

Information plays multiple roles in interaction between agents (be it humans or artificial entities, e.g., software agents). First, it can be the commodity that the agents compete for; in that case, it often defines the outcome of the “interaction game”. Key exchange protocols are a good example here, as the involved honest parties strive to learn the key of the other agent while at the same time preventing any information leak to the intruder. Secondly, information can define the semantic content of an action: typically, most actions specified in a security protocol consist in transmitting or processing some information. Thirdly, information can be a resource that enables actions and influences the outcome of the game. This is because agents need information to construct and execute plans that can be used to achieve their goals.

Most approaches to information flow security adopt the first perspective. That is, information defines the ultimate goal of the interaction. Classical information security properties specify *what* information must not leak, and *how* it could possibly leak (i.e., what channels of information leakage are considered), but they do not give account of *why* the information should not leak to the intruder. For example, the property of *noninterference* [18] assumes that the “low clearance” users cannot learn anything about the activities of the “high clearance” users. In order to violate this, the “low” users can try to analyse their observations and/or execute a sequence of explorative actions of their own.

Nondeducibility on strategies [54] makes the same assumption about *what* should not leak, but takes also into account covert channels that some “high” users can use to send signals to the “low” agents according to a previously agreed code. *Anonymity* in voting [10,16] captures that an observer cannot learn what candidate a particular has voted for by looking at the voter’s behavior, scanning the web bulletin board, coercing the voter to hand in the vote receipt, etc.

As a consequence, the classical properties of information security can only distinguish between relevant and irrelevant information leaks if the distinction is given explicitly as a parameter, e.g., by classifying available actions into sensitive and insensitive [18]. However, it is usually hard (if not impossible) to obtain such a distinction based on the internal characteristics of the actions. We illustrate the point below by means of a real-life example.

1.1 Motivating Example: Phone Banking

In some phone banking services, the maiden name of the user’s mother is used as a part of authentication, e.g., to change the settings of the account. That is, the user is typically asked to spell her name, birth-date, current address, and her mother’s maiden name in order to change the credit limit in the account, block/unblock ATM use in specified geographical areas, and so on. Note that information about one’s birth-date and address is fairly easy to obtain in public directories and/or repositories kept and marketed by various web services that require the data for registration. So, the mother’s maiden name plays the role of a “strong test of identity” in this scenario.³ Consider now a user posting an essay about some ancestor of hers on her blog, mentioning also the name of the ancestor. If the essay is about the user’s mother, it reveals potentially dangerous information. This is, among other things, because an intruder can use the information to: (1) access the phone banking service, (2) authenticate impersonating the user, (3) change (in the user’s profile) the telephone number used for web banking password recovery and sms authentication of web banking transactions, (4) change the web banking password of the user, and finally (5) log in and transfer money from the user’s account.

On the other hand, if the post is about some other member of the user’s family (father, grandmother, paternal grandfather, etc.) revealing the name of the person is probably harmless. Note that it is impossible to distinguish between the two pieces of information (say, the mother’s maiden name vs. the grandmother’s maiden name) based on their internal features. Both have the same syntactic structure of a single word (i.e., a string of characters with no blank spaces) and the same semantic content (a family name of a person; more precisely, the family name of the person at birth). The only difference lies in the context: the first kind of information is used in some important social procedures, while the second one is not.

³ This is a real-life example from the authors’ personal experience with BNP Paribas in one of EU countries. For similar security questions, used by various phone or web services, cf. e.g. [29].

1.2 Information as Strategic Resource

In this paper, we claim that a broader perspective is needed to appropriately model and analyse such scenarios. Agents compete for information not for its own sake, but for reasons that go beyond purely epistemic advantages. An intruder may want to know the password of a PayPal user in order to impersonate the user and steal some *real* money by making a payment to his own benefit (possibly via an account of a suitable “mule”). An industry player may need encryption keys used in internal communication between employees of its main competitor in order to find out about the competitor’s current business strategy. A political activist needs the ability to learn the value of someone’s vote in order to effectively coerce that person into voting for the candidate that the activist is rallying for. Thus, in most security scenarios, information is a resource rather than a commodity. More precisely, information is a commodity that the players compete for in an “information security game” but the game is played in the context of a “real” game where information is only a resource, enabling (some) players to achieve their non-epistemic goals. As players obtain new information, their uncertainty is reduced, and they increase their ability to choose a good strategy in the real game.

What would a *significant information leak* be in this view? To answer the question, we draw inspiration from the concept of the *value of information* from decision theory [25]: a piece of information is worth as much as it increases the expected payoff of the player. Similarly, an information leak is significant if it increases the ability of the attacker to construct a damaging attack strategy in the real game.

1.3 Main Idea and Contribution of the Paper

The main idea behind this paper can be summarized as follows. We consider three research questions:

- How can we evaluate the ability of an adversary to harm the goal of the system?
- How can we compare two systems with regard to the ability of attackers to harm the goal of the system in them?
- How can we know whether the ability (or inability) of the attacker to harm the goal of the system is because of some leakage of information to the attacker or not?

The paper is structured to discuss these questions in order. First, we use the concept of *surely winning strategies* from game theory to analyze the adversary’s strategic ability to disrupt the correct behavior of the system. This can be a functionality property, a security property, or a combination of the two kinds. Also, it can arise from a goal of the “high clearance” agents or from an objective assigned to the system by its designers and/or owners. Preventing the attacker from having a winning attack strategy is what the designer of the system may

want to achieve. We will see the *effective security* of the system as the attacker’s inability to come up with such a strategy.

Secondly, we use the notion of *effective security* for comparing two systems by looking at the strategic ability of an adversary to harm the goal of the system.

Thirdly, a successful attack strategy can exist due to flawed design of either the control flow or the information flow in the system. Here, we are interested in the latter. That is, we want to distinguish between vulnerabilities coming from the control vs. the information flow, and single out systems where redesigning the flow of information alone can make the system more secure. To this end, we define the noninterferent idealized variant of the system, which has the same control flow as the original system, but with the information reduced so that the system satisfies noninterference. Then, we define the system to be *effectively information-secure* if it is as good as its noninterfering idealized variant. As the main technical result, we show that the concept is well defined, i.e., the maximal noninterferent variant exists for every state-transition model.

We begin by presenting the preliminary concepts (models of interaction, noninterference, strategies and their outcomes) in Section 3. In Section 4, we define the generic concept of effective security. In Section 5, we look specifically at the security of information flow, and show how it can be defined based on the relation between the attacker’s observational capabilities and his ability to compromise the goals of the system. In Section 6 we extend our results to models that are not total on input. Finally, we summarize the work in Section 7.

2 Related Work

Various formalizations of information flow security have been proposed and studied. The classical concept here is *noninterference* [18] and its variations: *nondeducibility* [49], *noninference* [37], *restrictiveness* [33], *nondeducibility on strategies* [54], and *strategic noninterference* [26]. Although noninterference was originally introduced for finite transition systems, it was later redefined, generalized, and extended in the framework of process algebras [3,43,41,42,45]. Noninterference and its variants have been studied from different perspectives. Some works dealt with composability of noninterference [33,56,47]. Another group of papers studied the properties of intransitive noninterference [42,7,51,13] which is important in systems with downgraders. Probabilistic noninterference and quantitative noninterference have been investigated, e.g., in [21,54,34,38,31,48]. All the above concepts assume that the information flow in the system is secure only when no information ever flows from High to Low players. In this paper, we want to discard irrelevant information leaks, and only look at the significant ones (in the sense that the leaking information can be used to construct an attack on a higher-order correctness property).

The problem of how to weaken noninterference to successfully capture security guarantees for real systems has been also extensively studied. Most notably, postulates and policies for *declassification* (called also *information release*) were studied, cf. [46] for an introduction. This submission can be viewed as an at-

tempt to determine *what information is acceptable to declassify*. In this sense, our results can be useful in proposing new declassification policies and evaluating existing ones. We note, however, that the existing work on declassification are mainly concerned by the question *what* information can be released, *when*, *where*, and by *whom*. In contrast, we propose an argument for *why* it can be released. Moreover, declassification is typically about intentional release of information, whereas we do not distinguish between intentional and accidental information flow. Finally, the research on declassification assumes that security is defined by some given “secrets” to be protected. In our approach, no information is intrinsically secret, but the information flow is harmful if it enables the attacker to gain more strategic ability against the goals of the system.

Parameterized noninterference [17] can be seen as a theoretical counterpart of declassification, where security of information flow is parameterized by the analytic capabilities of the attacker. Again, that research does not answer why some information must be kept secret while some other needs not, and in particular it does not take strategic power of the attacker into account.

Economic and strategic analysis of security properties is a growing field in general, cf. e.g. [6,36,12,9,55,27]. A number of papers have applied game-theoretic concepts to define the security of information flow [32,23,24,11,15,26]. However, most of those papers [32,23,24,11] use games only in a narrow mathematical sense to provide a proof system (called the *game semantics*) for deciding security properties. We are aware of only a handful of papers that investigate the impact of participants’ incentives and available strategies on the security of information flow. In [2,22], economic interpretations of privacy-preserving behavior are proposed. [15] uses game-theoretic solution concepts (in particular, Nash equilibrium) to prescribe the optimal defense strategy against attacks on information security. In contrast, our approach is analytic rather than prescriptive, as we do not propose how to manage information security. Moreover, in our view, privacy is not the goal but rather the means to achieve some higher-level objectives. Finally, [26] proposes a weaker variant of noninterference by allowing the High players to select an appropriate strategy, while here we look at the potential damage inflicted by adverse strategies of the Low users.

Our idea of looking at the unique most precise non-interfering variant of the system is related on the technical level to [17]. There, attackers displaying different analytical capabilities are defined by abstract interpretation, which leads to a lattice of noninterference variants with various strength. Attackers with weakened observational powers were also studied in [57].

3 Preliminaries

We begin by presenting the main ingredients that we are going to use in our proposal. First, we introduce simple models of interaction that slightly extend the classical approach of Goguen and Meseguer. Then we recall Goguen and Meseguer’s definition of noninterference that captures the property of secure information flow from the “insider” agents to the “outsiders”. Finally, we present

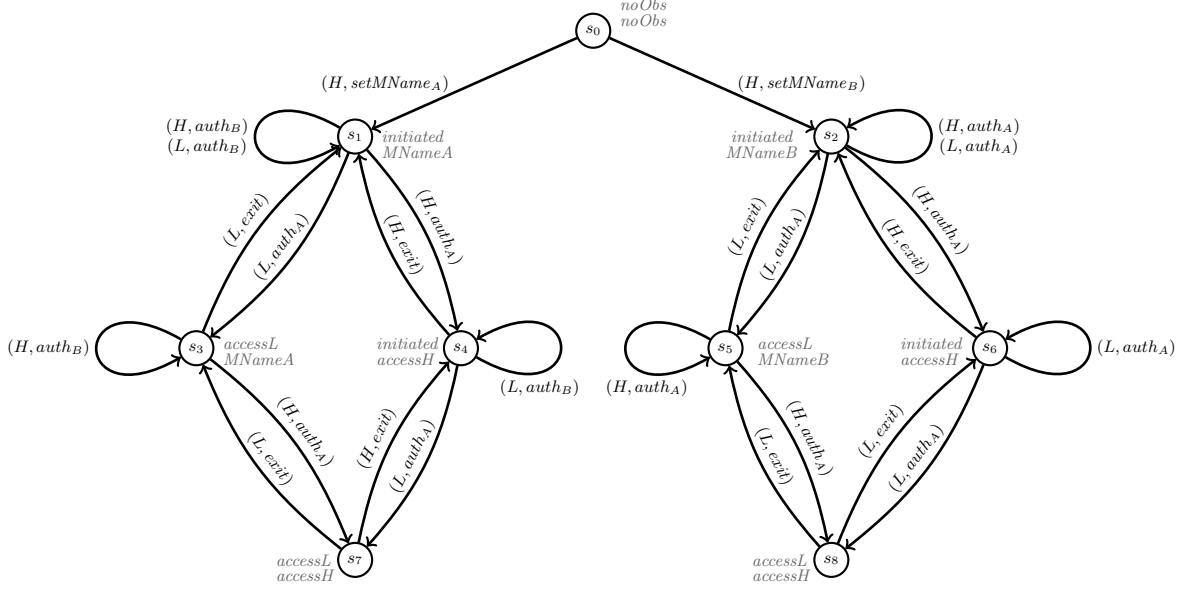


Fig. 1. A simple model for the phone banking example

some basic concepts from game theory (strategies, winning strategies) and theory of temporal specification (temporal goals).

3.1 Simple Models of Interaction

Since we build our proposal around the standard notion of noninterference by Goguen and Meseguer [18], we will use similar models to represent interaction between actions of different agents. The *system* is modeled by a multi-agent asynchronous transition network $M = \langle St, s_0, \mathcal{U}, \mathcal{A}, Obs, obs, do \rangle$ where: St is the set of *states*, s_0 is the initial state, \mathcal{U} is the set of *agents* (or *users*), \mathcal{A} is the set of *actions* (or *commands*), Obs is the set of possible *observations* (or *outputs*); $obs : St \times \mathcal{U} \rightarrow Obs$ is the observation function. $do : St \times \mathcal{U} \times \mathcal{A} \rightarrow St$ is the transition function that specifies the (deterministic) outcome $do(s, u, a)$ of action a executed by user u in state s . We will sometimes write $[s]_u$ instead of $obs(s, u)$. Also, we will call a pair $(user, action)$ a *personalized action*. We construct the multi-step transition function $exec : St \times (\mathcal{U} \times \mathcal{A})^* \rightarrow St$ so that, for a finite string $\alpha \in (\mathcal{U} \times \mathcal{A})^*$ of personalized actions, $exec(s, \alpha)$ denotes the state resulting from execution of α from s on. We may sometimes write $s \xrightarrow{\alpha} t$ instead of $exec(s, \alpha) = t$, and $exec(\alpha)$ instead of $exec(s_0, \alpha)$. The way models are constructed is illustrated by the following example.

Example 1. Consider a simplified version of the phone banking scenario from Section 1.1, where a client can access his/her account by correctly giving the

maiden name of his/her mother. Figure 1 presents the simplest possible transition network for the scenario. Labels on transitions show the personalized actions resulting in the transition, and the observations of users in each state are shown beside the state. There are two users: H who has an account in the bank, and L who may try to impersonate H . At the initial state, H enters her mother's maiden name when setting up her profile at the bank. To keep the graph simple, we include only two possibilities: action $setMName_A$ fixes the name as "A", whereas action $setMName_B$ sets the entry to "B". Clearly, H can observe which value she entered (different observations in states s_1, s_2). Moreover, L cannot observe that, as L 's observations in s_1, s_2 are the same.

Any user can access H 's bank account by mentioning the name correctly. So, depending on the value that has been entered, doing either $auth_A$ or $auth_B$ will grant the user with access to the bank account. Moreover, the user who authenticated successfully can observe it, and the other user cannot even see that the authentication took place. The user who has successfully logged in to the account, can log out by executing the action $exit$. To make the graph simpler, we also assume that giving the wrong name has no effect on the state of the system.

Three remarks are in order. First, Goguen and Meseguer's models define agents' observations based on states only, whereas it is often convenient to also model the information flow due to observing each others' actions. Secondly, the models are fully asynchronous in the sense that if each user "submits" a sequence of actions to be executed then every interleaving of the submitted sequences can occur as the resulting behavior of the system. No synchronization is possible. Thirdly, the models are "total on input" (each action label is available to every user at every state), and hence no synchronization mechanism can be encoded via availability of actions. Especially the last two features imply that models of Goguen and Meseguer allow for representation of a very limited class of systems.

More expressive classes of models include various kinds of transition systems [53], concurrent programs [28], interpreted systems [14], reactive modules [4], multi-agent transition networks (a.k.a. concurrent game structures) [5], and many more.

We start by using the purely asynchronous models of Goguen and Meseguer. Then, in Section 6, we extend our results to a broader class of models by allowing partial transition functions.

3.2 Noninterference

We now recall the standard notion of noninterference from [18]. Let $U \subseteq \mathfrak{U}$ and $\alpha \in (\mathfrak{U} \times \mathfrak{A})^*$. By $Purge_U(\alpha)$ we mean the subsequence of α obtained by eliminating all the pairs (u, a) with $u \in U$.

Definition 1 (Noninterference [18]). *Let M be a transition network with sets of "high clearance" agents H and "low clearance" agents L , such that $H \cap L = \emptyset$, $H \cup L = \mathfrak{U}$. We say that H is non-interfering with L iff for all $\alpha \in (\mathfrak{U} \times \mathfrak{A})^*$*

Example 2. Consider a simplified version of the phone banking scenario from Section 1.1. There are two users: H who has an account in the bank, and L who may try to impersonate H . H can access her account by correctly giving the maiden name of her mother. Moreover, H runs a blog, and can publish some of her personal information on it. We consider two alternative variants: one where H publishes her grandmother’s maiden name on the blog (Figure 2), and one where she publishes her mother’s maiden name (Figure 3). We assume that the possible names are A and B in the former case, and C and D in the latter. The observations of L are shown beside each state. The observations for H are omitted, as they will be irrelevant for our analysis.

Each model begins by initialization of the relevant names, represented by virtual actions of agent H . Then, H publishes an essay on her blog. In the first variant, the essay mentions the maiden name of H ’s grandmother. In the second variant, it mentions her mother’s maiden name. After H has published the essay, L can check the blog (action $chkWeb$). The resulting observation of L depends on what is published. Then, authentication proceeds like in Example 1: in order to log in, a user must give the correct value of H ’s mother’s maiden name.

Note that, for mathematical completeness, we must define the outcome of every user-action pair in every state. We assume that there are two “error states” s_{HErr} , s_{LErr} in models M_a and M_b (not shown in the graphs). Any action of H not depicted in the figure leads to s_{HErr} , and any action of L not depicted in the figure leads to s_{LErr} . We will later use the error states in the definition of the players’ goals, in such a way that L will always want to avoid s_{LErr} and H will want to avoid s_{HErr} . This way we can (however imperfectly) simulate some synchronization in the restricted framework of Goguen and Meseguer.

Neither M_a nor M_b satisfies noninterference from H to L . For instance, in the model of Figure 2, if $\alpha = \langle (H, setMName_A), (H, setGName_D), (H, publish), (L, chkWeb) \rangle$, the observation of L after sequence α is $GNameD$, but the observation of L after $Purge_H(\alpha) = \langle (L, chkWeb) \rangle$ is $noObs$, which is clearly different. \square

Again, two remarks are in order. First, noninterference focuses solely on the information flow in the system. If L can detect any activity of H then noninterference is lost, regardless of the nature of the activity and the possible uses of the information. In real systems, the impact of information flow goes well beyond the information itself. Information is sought and preserved for a reason, not for its own sake. Typically, L want to obtain information about H because they want to use it to achieve their goals more effectively (i.e., conclude a business contract, submit a better bid in an auction, get unauthorised access to a bank account etc.). On the other hand, H want to protect their private information from L because their goals may be in conflict with the goals of L . This is especially the case when the Low players are labeled as “attackers” or “intruders”.

Secondly, detecting H ’s actions may require L to engage in “diagnostic” activity, i.e., executing a sequence of actions whose only purpose is to determine if H was active or not. This becomes an issue when we see information as a resource used to obtain one’s goals, rather than *the* goal of the user’s activity. Then,

obtaining more information about H can be in conflict with what L must do in order to achieve their real goals. Thus, on one hand L need more information to construct a better strategy for their goals, but on the other hand to acquire the information they may have to depart from the successful strategy.

3.3 Strategies and Their Outcomes

Strategy is a game-theoretic concept which captures behavioral policies that an agent can consciously follow in order to realize some objective [52,30]. We begin with an abstract formulation, and mention the most representative examples of strategy types in the next paragraph. Let $T(M)$ be the *tree unfolding* of M . Also if $U \subseteq \mathfrak{U}$ is a subset of agents, let T' be a U -trimming of tree T iff T' is a subtree of T starting from the same root and obtained by removing an arbitrary subset of transitions labeled by actions of agents from U . For the moment, we assume that each subset of agents $U \subseteq \mathfrak{U}$ is assigned a set of available coalitional strategies Σ_U . The most important feature of a strategy $\sigma_U \in \Sigma_U$ is that *it constrains the possible behaviors of the system*. We represent it formally by the *outcome function* $out_M(\sigma_U)$ that removes the executions of the system that strategy σ_U would never choose. Therefore, for every $\sigma_U \in \Sigma_U$, its outcome $out_M(\sigma_U)$ is a U -trimming of $T(M)$.

Let h be a node in tree T corresponding to a particular finite history of interaction. We denote the sequence of personalized actions leading to h by $act^*(h)$. Furthermore, $act^*(T) = \{act^*(h) \mid h \in nodes(T)\}$ is the set of finite sequences of personalized actions that can occur in T .

Types of strategies. Strategies are usually constructed as mappings from possible situations that the player can recognize in the game, to actions of the player (or subsets of actions if we allow for nondeterministic strategies). Two types of such strategies are commonly used in the literature on game-like interaction: positional strategies and perfect recall strategies. Positional strategies represent conditional plans where the decision is solely based on what the agents see in the current state of the system, while perfect recall strategies capture conditional plans where the agents can base their decisions on the whole history of the game until that moment.

Positional strategies represent conditional plans where the decision is solely based on what the agents see in the current state of the system. Formally, for $u \in \mathfrak{U}$, the set of individual positional strategies of u is $\Sigma_u^{pos} = \{\sigma_u : St \rightarrow \mathcal{P}(\mathfrak{A}) \setminus \{\emptyset\} \mid \forall q, q' \in St \cdot [q]_u = [q']_u \Rightarrow \sigma_u(q) = \sigma_u(q')\}$, where $\mathcal{P}(X)$ denotes the powerset of X . Notice the “uniformity” constraint which enforces that the agent must specify the same action(s) in states with the same observations. Now, coalitional positional strategies for a group of agents $U \subseteq \mathfrak{U}$ are simply tuples of individual strategies, i.e., $\Sigma_U^{pos} = \times_{u \in U} (\Sigma_u^{pos})$. The outcome of $\sigma_U \in \Sigma_U^{pos}$ in model M is the tree obtained from $T(M)$ by removing all the branches that begin from a node containing state q with a personalized action $(u, a) \in U \times \mathfrak{A}$ such that $a \notin \sigma_U(q)$.

In this work we focus on adversaries playing perfect recall strategies.

Perfect recall strategies. Formally, the set of *perfect recall strategies* of agent u is $\Sigma_u^{\text{rec}} = \{\sigma_u : \text{nodes}(T(M)) \rightarrow \mathcal{P}(\mathfrak{A}) \setminus \{\emptyset\} \mid \text{obs}_u(h) = \text{obs}_u(h') \Rightarrow \sigma_u(h) = \sigma_u(h')\}$, where $\text{obs}_u(h)$ denotes the accumulate observations collected by agent u along history h . How to define obs_u for sequences of states? For asynchronous systems, this is typically defined as $\text{obs}_u(q) = [q]_u$, $\text{obs}_u(h \circ q) = \text{obs}_u(h)$ if $\text{last}(h) = q$, and $\text{obs}_u(h \circ q) = \text{obs}_u(h) \circ [q]_u$ otherwise (where \circ denotes the concatenation operator). That is, what u has learned along h is equivalent to the sequence of observations she has seen, modulo removal of “stuttering” observations. Now, coalitional strategies of perfect recall for a group of agents $U \subseteq \mathfrak{A}$ are combinations of individual strategies, i.e., $\Sigma_U^{\text{rec}} = \times_{u \in U} (\Sigma_u^{\text{rec}})$. The outcome of $\sigma_U \in \Sigma_U^{\text{rec}}$ in model M is the tree obtained from $T(M)$ by removing all the branches that begin from a node h with a personalized action $(u, a) \in U \times \mathfrak{A}$ such that $a \notin \sigma_U(h)$.

3.4 Temporal Goals and Winning Strategies

A goal is a property that some agents may attempt to enforce by selecting their behavior accordingly. In game-theoretic models, goals are typically phrased as properties of the final state in the game. In our case, there is no final state – the interaction can go on forever. Because of that, we understand goals as properties of the full temporal trace that executes the sequence of actions selected by users. We base our approach on the concepts of *paths* and *path properties*, used in temporal specification and verification of systems [8,35]. Let $\text{paths}(M)$ denote the set of infinite sequences of states that can be obtained by subsequent transitions in M . Additionally, we will use $\text{paths}_M(\sigma)$ as a shorthand for $\text{paths}(\text{out}_M(\sigma))$.

Definition 2 (Temporal goal [35]). A goal in M is any $\Gamma \subseteq \text{paths}(M)$. Note that $\text{paths}(M) = \text{paths}(T(M))$, so a goal can be equivalently seen as a subset of paths in the tree unfolding of M .

Most common examples of such goals are safety and reachability goals.

Definition 3 (Safety and reachability goals [35]). Given a set of safe states $\mathbb{S} \subseteq \text{St}$, the safety goal $\Gamma_{\mathbb{S}}$ is defined as $\Gamma_{\mathbb{S}} = \{\lambda \in \text{paths}(M) \mid \forall i. \lambda[i] \in \mathbb{S}\}$. Moreover, given a set of target states $\mathbb{T} \subseteq \text{St}$, the reachability goal $\Gamma_{\mathbb{T}}$ can be defined as $\Gamma_{\mathbb{T}} = \{\lambda \in \text{paths}(M) \mid \exists i. \lambda[i] \in \mathbb{T}\}$.

Definition 4 (Winning strategies). Given a transition network M , a set of agents $U \subseteq \mathfrak{A}$ with goal Γ_U , and a set of strategies Σ_U^{rec} , we say that U have a (surely winning) strategy to achieve Γ_U iff there exists a strategy $\sigma_U \in \Sigma_U^{\text{rec}}$ such that $\text{paths}_M(\sigma_U) \subseteq \Gamma_U$.

Example 3. Consider the models in Figure 2 and Figure 3, and suppose that L wants to access H ’s bank account. This can be expressed by the reachability goal $\Gamma_{\mathbb{T}}$ with $\mathbb{T} = \{s_{15}, s_{16}\}$ as the target states. In fact, L also wins if H executes an out-of-place action (cf. Example 2 for detailed explanation). In consequence, the winning states for L are $\mathbb{T} = \{s_{15}, s_{16}, s_{\text{HErr}}\}$. Note that L has no strategy

that guarantees Γ_T in model M_a (although information is theoretically leaking to L as the model does not satisfy noninterference). Even performing the action *chkWeb* does not help, because L cannot distinguish between states s_{11} and s_{13} , and there is no single action that succeeds for both s_{11}, s_{13} . Thus, L does not know whether to use *auth_A* or *auth_B* to get access to H 's bank account.

On the other hand, L has a winning strategy for Γ_T in model M_b . The strategy is to execute *chkWeb* after H publishes her mother's maiden name, and afterwards do *auth_A* in states s_{11}, s_{12} (after observing *MNameA*) or *auth_B* if the system gets to s_{13}, s_{14} (i.e., after observing *MNameB*). \square

In what follows, we will look at the L 's strategic ability to harm desirable behavior of the system.

4 Security as Strategic Property

The property of noninterference looks for *any* leakage of *any* information. If one can possibly happen in the system, then the system is deemed insecure. In many cases, this view is too strong. There are lots of information pieces that can leak out without bothering any interested party. Revealing the password to your web banking account can clearly have much more disastrous effects than revealing the price that you paid for metro tickets on your latest trip to Paris. Moreover, the relevance of an information leak cannot in general be determined by the type of the information. Think, again, of revealing the maiden name of your mother vs. the maiden name of your grandmother. The former case is potentially dangerous since the maiden name of one's mother is often used to grant access to manage banking services by telephone. Revealing the latter is quite harmless to most ends and purposes.⁴

In this paper, we suggest that the relevance of information leakage should be judged by the extent of damage that the leak allows the attackers to inflict on the goal of the system. Thus, as the first step, we define the security of the system in terms of damaging abilities of the Low players.

In order to assess the relevance of information flow from High to Low, we will look at the resulting strategic abilities of Low. For this, two design choices have to be made. First, what type of strategies are Low supposed to use? Secondly, what is the goal that they are assumed to pursue? The second question is especially important, because typically we do not know (and often do not care about) the real goals of potential attackers. What we know, and what we want to protect, is the objective that the system is built for.

We follow the game-theoretic tradition of looking at the worst case and assuming the opponents to be powerful and adversary. Thus, we assume L to use perfect recall strategies. Moreover, we assume that the goal of L is to violate a given goal of the system. The goal can be a functionality or a security requirement, or a combination of both. Moreover, it can originate from a private goal

⁴ Note, however, that revealing the maiden name of your maternal grandmother is potentially dangerous to your *mother* if she enables banking by telephone.

of the High players, an objective ascribed to the system by its designer (e.g., the designer of a contract signing protocol), or a combination of requirements specified by the owner/main stakeholder in the system (for instance, a bank in case of a web banking infrastructure).

Definition 5 (Effective security). *Let M be a transition network with some Low players $L \subseteq \mathfrak{U}$, and let Γ be the goal of the system. We say that M is effectively secure for (L, Γ) iff L does not have a strategy to enforce $\overline{\Gamma}$, where \overline{X} denotes the complement of set X . That is, the system is effectively secure iff the attackers do not have a strategy that ensures an execution violating the goal of the system. We will use $ES(M, L, \Gamma)$ to refer to this property.*

Besides judging the effective security of a system, we can also use the concept to compare the security level of two models.

Definition 6 (Comparative effective security). *Let M, M' be two models, and Γ be a goal in M, M' (i.e., $\Gamma \subseteq \text{paths}(M) \cup \text{paths}(M')$). We say that:*

- M has strictly less effective security than M' for (L, Γ) , denoted $M \prec_{L, \Gamma} M'$, iff $ES(M', L, \Gamma)$ but not $ES(M, L, \Gamma)$. That is, L can enforce a behavior of the system that violates its goal in model M but not in M' . We denote the relationship by $M \prec_{L, \Gamma} M'$;
- M' is at least as effectively secure as M for (L, Γ) , denoted $M \preceq_{L, \Gamma} M'$, iff $ES(M, L, \Gamma)$ implies $ES(M', L, \Gamma)$;
- M is effectively equivalent to M' for (L, Γ) , denoted $M \simeq_{L, \Gamma} M'$, iff either both $ES(M, L, \Gamma)$ and $ES(M', L, \Gamma)$ hold, or both do not hold.

Thus, if in one of the models L can construct a more harmful strategy then the model displays lower effective security than the other model. Conversely, if both models allow only for the same extent of damage then they have the same level of effective security. This way, we can order different alternative designs of the system according to the strategic power they give away to the attacker.

Example 4. Consider models M_a, M_b from Figure 2 and Figure 3, and let the goal Γ be to prevent L from accessing H 's bank account. Thus, Γ is the safety goal $\Gamma_{\mathbb{S}}$ with $\mathbb{S} = St \setminus \{s_{15}, s_{16}, s_{HErr}\}$, and therefore $\overline{\Gamma} = \Gamma_{\mathbb{T}}$ with $\mathbb{T} = \{s_{15}, s_{16}, s_{HErr}\}$. As we saw in Example 3, L has no strategy to guarantee $\overline{\Gamma}$ in M_a , but she has a surely winning strategy for $\overline{\Gamma}$ in M_b . Thus, M_b is strictly less effectively secure than M_a , i.e., $M_b \prec_{L, \Gamma} M_a$.

We will further use the concept to compare security of alternative information flows based on the same (or similar) action-transition structures.

5 Effective Information Security

We will now propose a scheme that allows to determine whether a given model of interaction leaks relevant information or not. We use the idea of refinement

checking from process algebras, where a process is assumed correct if and only if it refines the ideal process [40]. A similar reasoning scheme is also used in analysis of multi-party computation protocols (a protocol is correct iff it is equivalent to the ideal model of the computation [20]).

To this end, we need a suitable notion of refinement or equivalence, and a suitable definition of the ideal model. The former is straightforward: we will use the $\simeq_{L,\Gamma}$ relation. The latter is more involved. If the reference model ascribes too much observational capabilities to the Low players then the concept will be ill-defined (it will classify insecure systems as secure). If the reference model assigns Low with too little information then the concept will be useless (no realistic system will be ever classified as secure).

In what follows, we first explain and define the concept of an idealized variant of a model. Then in Section 5.2 we do our first take on the idealized variant by defining the *blind variant* of a model. In Section 5.3 we first define the *non-interfering idealized* variant of a model.

5.1 Ability-Based Security of Information Flows

Definition 6 allows for comparing the effective security of two alternative information flows. We will say two models differ only in their information flow if they are *transition equivalent*:

Definition 7 (Transition-equivalent models). *The action-transition frame of a model M , which we denote by F_M , is the network M minus the observation functions $\text{obs}(\cdot)$. We will denote the set of models based on frame F by $\mathcal{M}(F)$. Two models are transition-equivalent iff they are based on the same frame.*

Then, $(F, \text{obs}) \prec (F, \text{obs}')$ says that the observation function obs leaks more relevant information than obs' in the transition-action frame F . However, we usually do not want to compare several alternative information flows. Rather, we want to determine if a single given model M reveals relevant information or not. How can we achieve that? A natural idea is to compare the effective security of M to an *ideal model*, i.e. a model that is transition equivalent to the original model and moreover leaks no relevant information by construction. Then, a model is effectively information-secure if it has the same level of effective security as its idealized variant:

Definition 8 (Effective information security). *Let M be a transition network with some Low players $L \subseteq \mathfrak{U}$, and let Γ be the goal of the system. Moreover, let $\text{Ideal}(M)$ be the idealized variant of M . We say that M is effectively information-secure for (L, Γ) iff $M \simeq_{L,\Gamma} \text{Ideal}(M)$.*

How do we construct the idealized variant of M ? The idea is to “blur” observations of Low so that we obtain a variant of the system where the observational capabilities of the attackers are minimal. What observational capabilities are “minimal”? We start with the following, rather naive definition of idealization.

5.2 Blinding the Low Players: First Attempt

By using the idealized model, we intend to distinguish to what extent the damaging abilities of Low are due to the “hard” actions available in the system, and to what extent they are due to the available information flow. In other words, we want to see how far one can minimize the strategic ability of the Low players by reducing their observational abilities in the model.

The first take to define an idealized model is to assume that L *never sees anything*. To this end, we simply assume that $obs(s, L)$ is the same in all states $s \in St$.

Definition 9 (Idealized model, first take). *Having a transition network M based on frame F , and a set of players L , we define the blind variant of M as $M' = (F, obs')$ such that $obs'(q, l) = obs(q', l)$ for every $q, q' \in St$ and $l \in L$.*

In many scenarios this is too much. In particular, a Low agent may have access to perfectly legitimate observations that are inherent to maintaining their private affairs, such as checking the balance of their bank account, listing the files stored on in their private file space, etc.

5.3 Idealized Models Based on Noninterference

Below we propose a weaker form of “blinding” that will be used to single out the damaging abilities that are due to Low *observing High’s actions*, rather than due to *any* observations that Low can happen to make. We begin by recalling the notion of *term unification* which is a fundamental concept in automated theorem proving and logic programming [39]. Given two terms t_1, t_2 , their unification ($t_1 \equiv t_2$) can be understood as a declaration that, from now on, both terms refer to exactly the same underlying object. In our case the terms are observation labels from the set Obs . A unification can be seen as an equivalence relation on observation labels, or equivalently as a partitioning of the labels into equivalence classes. The application of the unification to a model yields a similar model where the equivalent observations are “blurred”.

Definition 10 (Unification of observations). *Given a set of observation labels Obs , a unification on Obs is any equivalence relation $\mathcal{U} \subseteq Obs \times Obs$.*

Given a model $M = \langle St, s_0, \mathcal{A}, \mathcal{A}, do, Obs, obs \rangle$ and a unification $\mathcal{U} \subseteq Obs \times Obs$, the application of \mathcal{U} to M is the model $\mathcal{U}(M) = \langle St, s_0, \mathcal{A}, \mathcal{A}, do, Obs', obs' \rangle$, where: $Obs' = \{[o]_{\mathcal{U}} \mid o \in Obs\}$ replaces Obs by the set of equivalence classes defined by \mathcal{U} , and $obs'(q, u) = [obs(q, u)]_{\mathcal{U}}$ replaces the original observation in q with its equivalence class for any $u \in \mathcal{A}$.

Example 5. Figure 4 depicts the model obtained from M_b by unifying observations $MNameA$ and $MNameB$ into $\{MNameA, MNameB\}$, observations $init$ and $noObs$ into $\{init, noObs\}$, and observation $accessL$ into $\{accessL\}$.

Our reference model for M will be the variant of M where noninterference is obtained by the minimal necessary “blurring” of L ’s observations.

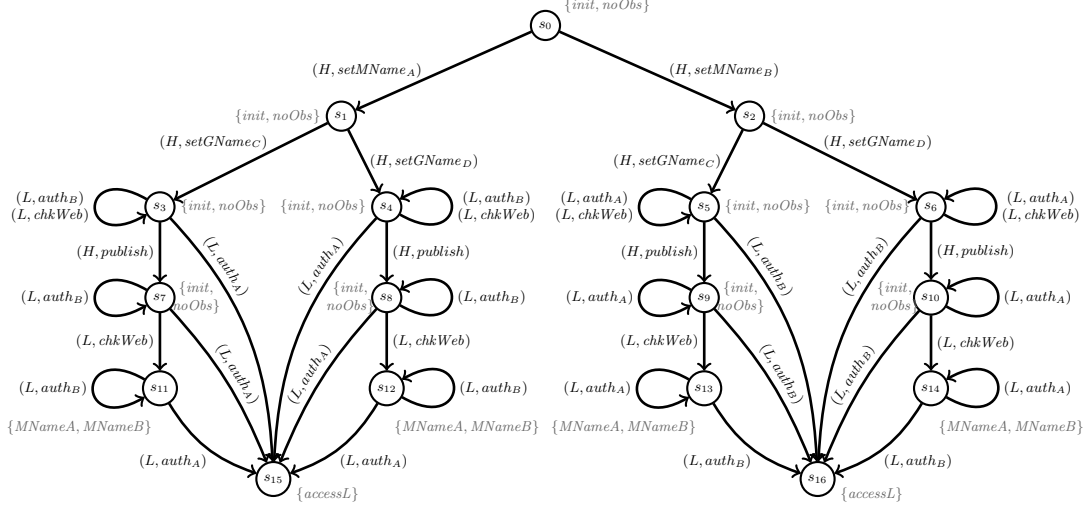


Fig. 4. An example of unification of observations.

Definition 11 (Noninterferent idealized model). *Having a transition network M and a set of “low” players L , we define the noninterfering idealized variant of M as $\mathcal{U}(M)$ such that:*

- (i) $NI_{\mathcal{U}(M)}(H, L)$, and
- (ii) for every $\mathcal{U}' \subsetneq \mathcal{U}$ it is not the case that $NI_{\mathcal{U}'}(H, L)$.

We need to show that the concept of noninterferent idealized model is well defined. The proof is constructive, i.e., given a model M , we first show how one can build its idealized variant, and then show that it is unique.

Theorem 1. *For every transition network M , there is always a unique unification \mathcal{U} satisfying properties (i) and (ii) from Definition 11.*

The proof of Theorem 1 needs some preliminary steps. As the first step, we recall and adapt the concept of unwinding relations [19,44,50]. Unwinding is constructed analogously to the standard notion of bisimulation, and requires Low’s uncertainty to be a fixpoint of an appropriate relation transformer. Unwinding relations are important because they characterize noninterference in purely structural terms. Moreover, existence of an unwinding relation is usually easier to verify than proving noninterference directly. We then use the concept of unwinding relation to define relation R_M^* on the states of a transition network M . We use this relation to construct and prove the uniqueness of the idealized variant of M .

Definition 12 (Unwinding for Noninterference [50]). *Let M be a transition network, H a set of High agents, and L a set of Low agents. Then,*

$\sim_{NI_L} \subseteq St \times St$ is an unwinding relation iff it is an equivalence relation satisfying the conditions of output consistency (OC), step consistency (SC), and local respect (LR). That is, for all states $s, t \in St$:

- (OC) If $s \sim_{NI_L} t$ then $[s]_L = [t]_L$;
- (SC) If $s \sim_{NI_L} t$, $u \in L$, and $a \in \mathfrak{A}$ then $do(s, u, a) \sim_{NI_L} do(t, u, a)$;
- (LR) If $u \in H$ and $a \in \mathfrak{A}$ then $s \sim_{NI_L} do(s, u, a)$.

Proposition 1 ([50]). $NI_M(H, L)$ iff there exist an unwinding relation \sim_{NI_L} on the states of M that satisfies (OC), (SC) and (LR).

Next we define R_M^* on the states of a transition network M . The definition goes as follows: first we relate any two states of M' if one of them can be reached from the other one by a sequence of High personalized actions. Then in each step we relate the pair of states that are reached by a similar Low personalized action from any two states that are already related. Also, we enforce transitivity on the set. We continue adding related states until the relation becomes stable. The mathematical definition of R_M^* is as follows:

Definition 13 (Relation R_M^* for a transition network M). Given a model $M = \langle St, s_0, \mathfrak{U}, \mathfrak{A}, do, Obs, obs \rangle$ and sets of High players H and Low players L , we define the relation $R_M^* \subseteq St \times St$ as the least fixpoint of the following function F , transforming relations on St :

$$F(R) = R_0 \cup \{ (t_1, t_2) \mid \exists (s_1, s_2) \in R, l \in L, a \in \mathfrak{A}. do(s_1, l, a) = t_1, do(s_2, l, a) = t_2 \} \cup \{ (t_1, t_2) \mid \exists s \in St. (t_1, s) \in R \& (s, t_2) \in R \},$$

where $(s_1, s_2) \in R_0$ iff for some sequence of personalized actions of High players α , either $s_1 \xrightarrow{\alpha} s_2$, or $s_2 \xrightarrow{\alpha} s_1$.

It is straightforward to see that R_M^* is an equivalence relation (for reflexivity, notice that for any $s \in St$, $s \xrightarrow{\alpha} s$ for $\alpha = \langle \rangle$, and therefore $(s, s) \in R_0$). We will now show that if M satisfies noninterference then R_M^* is the smallest unwinding relation. Conversely, if M does not satisfy noninterference then R_M^* indicates pairs of states that must bear the same observations for Low if we want to make the model M non-interferent. We will later show that it is sufficient to unify Low's observations in states connected by R_M^* in order to obtain a non-interferent variant of M . In consequence, R_M^* generates the minimal unification that achieves the task.

The following proposition shows that if M satisfies the noninterference property, then R_M^* is a subset of any unwinding relations on the states of M .

Proposition 2. Given a model M and sets of players H and L , if \sim_{NI_L} is an unwinding relation on the states of M and relation R_M^* is defined as in Definition 13, then $R_M^* \subseteq \sim_{NI_L}$.

Proof. See the Appendix.

Now we show that if the model M , L players have the same observations at any two states related by R_M^* , then M satisfies noninterference.

Lemma 1. *In a model M with sets of players H and L , if for any $l \in L$, $s_1, s_2 \in St$ it is the case that $(s_1, s_2) \in R_M^*$ implies $obs(s_1, l) = obs(s_2, l)$, then R^* is an unwinding relation on the states of M and therefore it holds that $NI_M(H, L)$.*

Proof. We prove this by showing that R_M^* satisfies the conditions of Definition 12: The relation R_M^* is an equivalence relation, condition (OC) follows from the assumption of this lemma, and conditions (SC) and (LR) follow from the definition of the relation R_M^* . Therefore it holds that $NI_M(H, L)$.

And as the last step before introducing the unification of function \mathcal{U}_M^* , we show that if M satisfies noninterference, then R_M^* is an unwinding relation on its states (and by Proposition 2 it is in fact the smallest unwinding relation).

Proposition 3. *In a model M with sets of players H and L , if $NI_M(H, L)$ then R_M^* is an unwinding relation on states of M .*

Proof. See the Appendix.

Now, by using relation R_M^* , we define the unification of observations \mathcal{U}_M^* that will provide the noninterferent idealized variant of M .

Definition 14 (Unification for noninterference \mathcal{U}_M^*). *We define the unification of observations $\mathcal{U}_M^* \subseteq Obs \times Obs$ as follows. For any $o_1, o_2 \in Obs$, we have $(o_1, o_2) \in \mathcal{U}_M^*$ iff there exist $s_1, s_2, t_1, t_2 \in St$ and $l \in L$ such that:*

- (a) $obs(s_1, l) = o_1$,
- (b) $obs(s_2, l) = o_2$,
- (c) $(s_1, t_1) \in R_M^*$,
- (d) $(s_2, t_2) \in R_M^*$, and
- (e) $obs(t_1, l) = obs(t_2, l)$.

It then holds that $\mathcal{U}_M^*(M)$ satisfies the noninterference property (Proposition 4) and no refinement of \mathcal{U}_M^* achieves that (Proposition 5). The following lemma states that if two states are related by R_M^* , then their observations are unified by \mathcal{U}_M^* .

Lemma 2. *In a model M , for any $s_1, s_2 \in St$ and $l \in L$, if $(s_1, s_2) \in R_M^*$ then $(obs(s_1, l), obs(s_2, l)) \in \mathcal{U}_M^*$.*

Proof. See the Appendix.

As the next step, we show that $\mathcal{U}_M^*(M)$ satisfies the noninterference property.

Proposition 4. *Given a model M , and $\mathcal{U}_M^*(M) = \langle St, s_0, \mathfrak{U}, \mathfrak{A}, do, Obs^*, obs^* \rangle$ defined as in Definition 14 on M , it holds that $NI_{\mathcal{U}_M^*(M)}(H, L)$.*

Proof. See the Appendix.

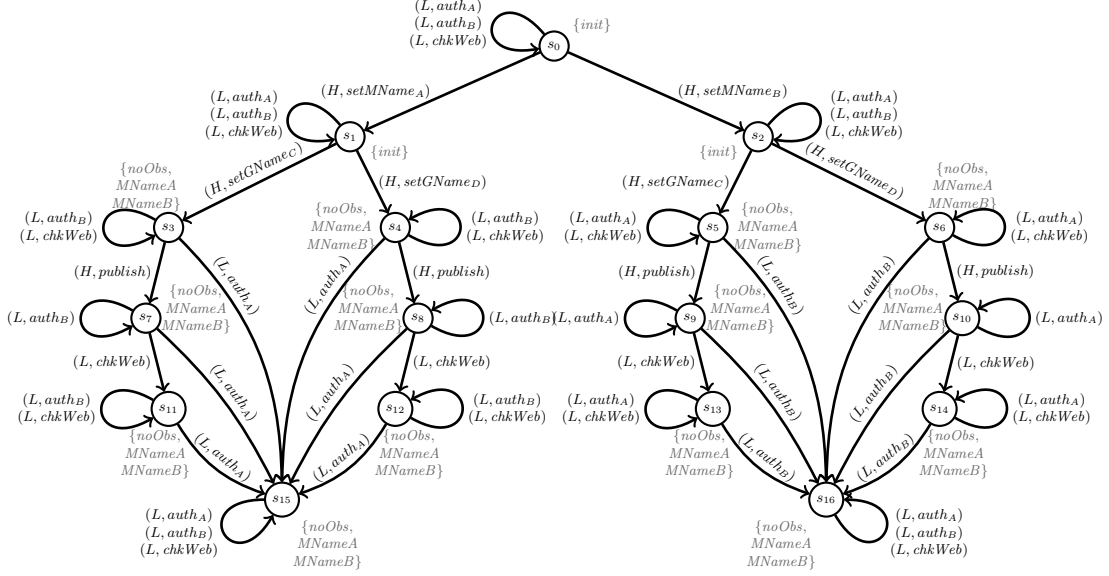


Fig. 5. The noninterfering idealized variant of the banking model M_b

As the last step before proving Theorem 1, we show that \mathcal{U}_M^* is the minimal unification that makes the model M noninterfering.

Proposition 5. *Given a model M , and sets of players H and L , for any unification of observations U where $U(M) = \langle St, s_0, \mathfrak{A}, \mathfrak{A}, do, Obs', obs' \rangle$, if $NI_{U(M)}(H, L)$ then $\mathcal{U}_M^* \subseteq U$.*

We can now complete the proof of Theorem 1.

Proof (of Theorem 1). We want to prove that, given a model M , set of players H and L , and any unification of observations \mathcal{U} , if $\mathcal{U}(M)$ is a noninterfering idealized variant of M , then $\mathcal{U} = \mathcal{U}_M^*$. Assume that $\mathcal{U}(M)$ is a noninterfering idealized variant of M . By property (i) of Definition 11 and Proposition 5 we infer that $\mathcal{U}_M^* \subseteq \mathcal{U}$. Also, by Proposition 4, we have that $NI_{\mathcal{U}_M^*(M)}(H, L)$. Therefore by property (ii) of Definition 11 it holds that $\mathcal{U} = \mathcal{U}_M^*$.

From now on, we assume that $Ideal(M)$ refers to the noninterfering idealized variant of M .

Example 6. Consider models M_a, M_b in Figure 2 and Figure 2. We recall that both models are not noninterferent. In the noninterferent idealized variant of M_a , observations $noObs$, $MnameC$, and $MNameD$ of L are unified and replaced by the equivalence class $\{noObs, MNameD, MNameD\}$. The idealized variant of M_b is constructed analogously by unification of $noObs$, $MnameA$,

and $MNameB$. Figure 5 shows the idealized variant $Ideal(M_b)$ of M_b . Clearly, L has no surely winning strategy to guarantee $\overline{I} = I_{\mathbb{T}}$ for $\mathbb{T} = \{s_{15}, s_{16}, s_{HErr}\}$ in both $Ideal(M_a)$ and $Ideal(M_b)$.

Recall from Example 4 that L has no winning strategy for \overline{I} in M_a , but she has one in M_b . So, $M_a \simeq_{L, \Gamma} Ideal(M_a)$, but $M_b \not\simeq_{L, \Gamma} Ideal(M_b)$. Thus, M_a is effectively information-secure for (L, Γ) , but M_b is not. \square

It is important to notice that noninterferent variants are indeed idealizations:

Proposition 6. *For every M , L , and Γ , we have that $M \preceq_{L, \Gamma} Ideal(M)$.*

Proof. Note that because M and $Ideal(M)$ differ only in their observation functions. Also we have that for any pair of states $s_1, s_2 \in St$, if $[s_1]_L^M = [s_2]_L^M$ then $[s_1]_L^{Ideal(M)} = [s_2]_L^{Ideal(M)}$. Therefore all the strategies of L in $Ideal(M)$ are also L 's strategies in M . Thus for any for any goal $\Gamma \subseteq paths(M)$, if L have a surely winning strategy to enforce \overline{I} in $Ideal(M)$ then they also have a surely winning strategy for \overline{I} in M , \square

Finally, note that the concept of noninterference in our construction of effective security can be in principle replaced by an arbitrary constraint of information leakage. The same reasoning scheme could be applied to noninference, nondeducibility, strategic noninterference, and so on. The pattern does not change: given a “classical” property \mathcal{P} of information security, we define the idealized variant of M through the minimal unification U such that that $U(M)$ satisfies \mathcal{P} . Then, M is effectively secure in the context of property \mathcal{P} iff it is strategically equivalent to $U(M)$.

We leave the investigation of which information security properties have unique minimal unifications for future work.

6 Extending the Results to a Broader Class of Models

As mentioned before, the models of Goguen and Meseguer are “total on input,” i.e., each action label is available to every user at every state. This makes modeling actual systems very cumbersome. We have seen that in the previous examples where spurious states had to be added to the analysis to allow for some synchronization between actions of different agents. In this section, we consider a broader class of models, and show how our results carry over to the more expressive setting. That is, we consider *partial transition networks (PTS)* $M = \langle St, s_0, \mathcal{U}, \mathcal{A}, Obs, obs, do \rangle$ which are defined as in Section 3.1, except that the transition function $do : St \times \mathcal{U} \times \mathcal{A} \rightarrow St$ can be a partial function. By $do(s, u, a) = undef$ we denote that action a is unavailable to user u in state s ; additionally, we define $act(s, u) = \{a \in \mathcal{A} \mid do(s, u, a) \neq undef\}$ as the set of actions available to u in s . Moreover, we assume that players are aware of their available actions, and hence can distinguish states with different repertoires of

choices – formally, for any $u \in \mathfrak{U}$, $s_1, s_2 \in St$, if $obs(s_1, u) = obs(s_2, u)$ then $act(s_1, u) = act(s_2, u)$.

We begin by a suitable update of the definition of noninterference:

Definition 15 (Noninterference for partial transition networks). *Given a PTS M and sets of agents H, L , such that $H \cup L = \mathfrak{U}$, $H \cap L = \emptyset$, we say that H is non-interfering with L iff for all $\alpha \in (\mathfrak{U} \times \mathfrak{A})^*$ and all $u_l \in L$, if $exec(\alpha) \neq undef$ then $[exec(\alpha)]_{u_l} = [exec(Purge_H(\alpha))]_{u_l}$. We denote the property also by $NI_M(H, L)$, thus slightly overloading the notation.*

Note that Definition 1 is a special case of Definition 15. We now define the noninterferent idealized variant based on the *total extension* of a PTS.

Definition 16 (U-total extension). *Given a PTS $M = \langle St, s_0, \mathfrak{U}, \mathfrak{A}, Obs, obs, do \rangle$ and a subset of users $U \subseteq \mathfrak{U}$, we define the U -total variant of M as $total_U(M) = \langle St, s_0, \mathfrak{U}, \mathfrak{A}, Obs, obs, do' \rangle$ where the transition function $do'(\cdot)$ is defined as follows: for every $s \in St$, $v \in \mathfrak{U}$ and $a \in \mathfrak{A}$, $do'(s, v, a) = s$ if for some $u \in U$ we have $v = u$ and $do(s, u, a) = undef$, otherwise $do'(s, v, a) = do(s, v, a)$.*

Definition 17 (Noninterferent idealized model for PTN). *Given a partial transition network M and a set of “low” players L , we define the noninterferent idealized variant of M as $\mathcal{U}(total_L(M))$ such that:*

- (i) $NI_{\mathcal{U}(total_L(M))}(H, L)$, and
- (ii) for every $\mathcal{U}' \subsetneq \mathcal{U}$ it is not the case that $NI_{\mathcal{U}'(total_L(M))}(H, L)$.

The uniqueness theorem is then stated similar to Theorem 1:

Theorem 2. *For every partial transition network M , there is always a unique unification \mathcal{U} satisfying properties (i) and (ii) from Definition 17.*

The proof is similar to the proof of Theorem 1, with the difference that we use $R_{total_L(M)}^*$ instead of R_M^* for constructing the idealized variant. However, as we use the concept of unwinding relation as the basis for using the R^* relation for constructing the idealized variant, we first need to modify the definition of the unwinding relation in Definition 12 and its corresponding proposition, Proposition 1 to adapt them to the new model:

Definition 18 (Unwinding for Noninterference in PTN). *Let M be a transition network, H a set of High agents, and L a set of Low agents. Then, $\sim_{NI_L} \subseteq St \times St$ is an unwinding relation iff it is an equivalence relation satisfying the conditions of output consistency (OC), step consistency (SC), and local respect (LR). That is, for all states $s, t \in St$:*

- (OC) If $s \sim_{NI_L} t$ then $[s]_L = [t]_L$;
- (SC) If $s \sim_{NI_L} t$, $u \in L$, and $a \in \mathfrak{A}$ then $a \in act(s, u)$ implies $do(s, u, a) \sim_{NI_L} do(t, u, a)$;
- (LR) If $u \in H$ and $a \in \mathfrak{A}$ then $a \in act(s, u)$ implies $s \sim_{NI_L} do(s, u, a)$.

Proposition 7. $NI_M(H, L)$ iff there exist an unwinding relation \sim_{NI_L} on the states of M that satisfies (OC), (SC) and (LR).

The rest of the proof of Theorem 2 follows analogously.

Example 7. With PTS, the scenario from Example 2 can be modeled directly, without spurious states that ruled out illegal transitions. Thus, our models M_a, M_b for the two variants of the scenario are now exactly depicted in Figures 2 and 3.

The noninterferent idealized variants of M_a (resp. M_b) is again obtained by the unification of observations $noObs$, $MnameC$, and $MNameD$ (resp. $noObs$, $MnameA$, and $MNameB$). Clearly, L has no surely winning strategy to guarantee $\overline{\Gamma} = \Gamma_{\mathbb{T}}$ for $\mathbb{T} = \{s_{15}, s_{16}\}$ in M_a , $Ideal(M_a)$, and $Ideal(M_b)$. Moreover, he has a surely winning strategy in M_b . In consequence, M_a is effectively information-secure for (L, F) , but M_b is not. \square

The noninterferent variant was indeed an idealization in simple transition networks of Goguen and Mesguer. Is it still the case in partial transition networks? That is, is it always the case that L has no more abilities in $Ideal(M)$ than in M ? In general, no. On one hand, L 's observational capabilities are more limited in $Ideal(M)$, and in consequence some strategies in M are no longer uniform in $Ideal(M)$. On the other hand, unification U^* possibly adds new transitions to M , that can be used by L in $Ideal(M)$ to construct new strategies. However, under some reasonable assumptions, $Ideal(M)$ does provide idealization, as shown in the two propositions below.

Proposition 8. *Let M be a PTN such that for every state s in M there is at least one player $u \notin L$ with $act(s, u) \neq \emptyset$. Then, for any Γ , we have that $M \preceq_{L, \Gamma} Ideal(M)$.*

Proposition 9. *For any PTN M and safety goal Γ , we have $M \preceq_{L, \Gamma} Ideal(M)$.*

7 Conclusions

In this paper, we introduce the novel concept of *effective information security*. The idea is aimed at assessing the relevance of information leakage in a system, based on how much the leakage enables an adversary to harm the correct behavior of the system. This contrasts with the common approach to information flow security where revealing any information is seen as being intrinsically harmful. We say that two information flows are *effectively equivalent* if the strategic ability of the adversary is similar in both of them. Moreover, one of them is *less effectively secure* than the other one if the amount of information leaked to the adversary in it increases the damaging ability of the adversary.

In order to determine how critical the information leakage is in a given system, we compare the damaging ability of the adversary to his ability in the idealized variant of the model. We define idealized models based on noninterference, and show that the construction is well defined. We prove this first for the deterministic, fully asynchronous transition networks of Goguen and Meseguer,

and then extend the results to structures that allow for a more flexible modeling of interaction. The construction includes an algorithm that computes the idealized variant of each model in polynomial time wrt the size of the model.

Note that the concept of noninterference in our construction of effective security can be in principle replaced by an arbitrary property of information flow. The same reasoning scheme could be applied to noninference, nondeducibility, strategic noninterference, and so on. The pattern does not change: given a property \mathcal{P} , we define the idealized variant of M through the minimal unification U such that $U(M)$ satisfies \mathcal{P} . Then, M is effectively information-secure in the context of property \mathcal{P} iff it is strategically equivalent to $U(M)$. We leave the investigation of which information security properties have unique minimal unifications for future work. Moreover, we are currently working on a more refined version of effective information security based on coalitional effectivity functions [1], in which the strategic ability of the adversary is not only compared at the initial state of the system, but across the whole state space.

References

1. J. Abdou and H. Keiding. *Effectivity Functions in Social Choice*. Springer, 1991.
2. A. Acquisti and J. Grossklags. Privacy attitudes and privacy behavior - losses, gains, and hyperbolic discounting. In *Economics of Information Security*, volume 12 of *Advances in Information Security*, pages 165–178. Springer, 2004.
3. P.G. Allen. A comparison of non-interference and non-deducibility using CSP. In *Proceedings of CSFW*, pages 43–54, 1991.
4. R. Alur and T. A. Henzinger. Reactive modules. *Formal Methods in System Design*, 15(1):7–48, 1999.
5. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
6. R. Anderson, T. Moore, S. Nagaraja, and A. Ozment. Incentives and information security. In *Algorithmic Game Theory*. 2007.
7. Michael Backes and Birgit Pfizmann. Intransitive non-interference for cryptographic purposes. In *Proceedings of S&P*, pages 140–152. IEEE, 2003.
8. J.R. Büchi. On a decision method in restricted second order arithmetic. In *Logic, Methodology and Philosophy of Science. Proc. 1960 Intern. Congr.*, pages 1–11. Stanford University Press, 1962.
9. Ahto Buldas and Triinu Mägi. Practical security analysis of e-voting systems. In *Proceedings of IWSEC*, volume 4752 of *Lecture Notes in Computer Science*, pages 320–335. Springer, 2007.
10. D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24:84–90, 1981.
11. A.S. Dimovski. Ensuring secure non-interference of programs by game semantics. In *Security and Trust Management*, pages 81–96. Springer, 2014.
12. Y. Dodis and T. Rabin. Cryptography and game theory. In *Algorithmic Game Theory*, chapter 8, pages 181–208. 2007.
13. Kai Engelhardt, Ron van der Meyden, and Chenyi Zhang. Intransitive noninterference in nondeterministic systems. In *Proceedings of CCS*, pages 869–880. ACM, 2012.

14. R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
15. A. Fielder, E. Panaousis, P. Malacaria, C. Hankin, and F. Smeraldi. Game theory meets information security management. *IFIP Advances in Information and Communication Technology*, 428:15–29, 2014.
16. A. Fujioka, T. Okamoto, and K. Ohta. A practical secret voting scheme for large scale elections. In *Proceedings of AUSCRYPT*, pages 244–251, 1992.
17. R. Giacobazzi and I. Mastroeni. Abstract non-interference: parameterizing non-interference by abstract interpretation. In *Proceedings of POPL*, pages 186–197. ACM, 2004.
18. Joseph A Goguen and José Meseguer. Security policies and security models. In *Proceedings of S&P*, pages 11–20. IEEE Computer Society, 1982.
19. Joseph A Goguen and José Meseguer. Unwinding and inference control. In *IEEE Symposium on Security and Privacy*, pages 75–75. IEEE Computer Society, 1984.
20. O. Goldreich, S. Micali, and A. Wigderson. How to play ANY mental game. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing STOC '87*, pages 218–229. ACM, 1987.
21. James W Gray III. Probabilistic interference. In *Proceedings of S&P*, pages 170–179. IEEE, 1990.
22. J. Grossklags, N. Christin, and J. Chuang. Secure or insure? A game-theoretic analysis of information security games. In *Proceedings of WWW*, pages 209–218. ACM, 2008.
23. C. Hankin, R. Nagarajan, and P. Sampath. Flow analysis: Games and nets. In *The Essence of Computation, Complexity, Analysis, Transformation. Essays Dedicated to Neil D. Jones [on occasion of his 60th birthday]*, volume 2566 of *Lecture Notes in Computer Science*, pages 135–156. Springer, 2002.
24. W.R. Harris, S. Jha, T.W. Reps, J. Anderson, and R.N.M. Watson. Declarative, temporal, and practical programming with capabilities. In *Proceedings of SP*, pages 18–32. IEEE Computer Society, 2013.
25. R. A. Howard. Information value theory. *IEEE Transactions on Systems Science and Cybernetics*, pages 22–26, 1966.
26. W. Jamroga and M. Tabatabaei. Strategic noninterference. In *Proceedings of the 30th International Conference on ICT Systems Security and Privacy Protection IFIP SEC 2015*, volume 455 of *IFIP Advances in Information and Communication Technology*, pages 67–81. Springer, 2015.
27. D. Korzhyk, Z. Yin, C. Kiekintveld, V. Conitzer, and M. Tambe. Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research*, 41:297–327, 2011.
28. O. Kupferman, M.Y. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking. *Journal of the ACM*, 47(2):312–360, 2000.
29. J. Levin. In what city did you honeymoon? and other monstrously stupid bank security questions. *Slate*, 2008.
30. K. Leyton-Brown and Y. Shoham. *Essentials of Game Theory: A Concise, Multi-disciplinary Introduction*. Morgan & Claypool, 2008.
31. Peng Li and Steve Zdancewic. Downgrading policies and relaxed noninterference. In *ACM SIGPLAN Notices*, volume 40, pages 158–170. ACM, 2005.
32. P. Malacaria and C. Hankin. Non-deterministic games and program analysis: An application to security. In *Proceedings of LICS*, pages 443–452. IEEE Computer Society, 1999.

33. Daryl McCullough. Noninterference and the composability of security properties. In *Proceedings of S&P*, pages 177–186. IEEE, 1988.
34. Annabelle McIver and Carroll Morgan. A probabilistic approach to information hiding. *Programming Methodology*, pages 441–460, 2003.
35. R. McNaughton. Testing and generating infinite sequences by a finite automaton. *Information and Control*, 9:521–530, 1966.
36. T. Moore and R. Anderson. Economics and internet security: a survey of recent analytical, empirical and behavioral research. Technical Report TR-03-11, Computer Science Group, Harvard University, 2011.
37. Colin O’Halloran. A calculus of information flow. In *Proceedings of ESORICS*, pages 147–159, 1990.
38. Alessandra Di Pierro, Chris Hankin, and Herbert Wiklicky. Approximate non-interference. *Journal of Computer Security*, 12(1):37–81, 2004.
39. J. A. Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, 1965.
40. A. W. Roscoe, C. A. R. Hoare, and R. Bird. *The Theory and Practice of Concurrency*. Prentice Hall PTR, 1997.
41. A.W. Roscoe. CSP and determinism in security modelling. In *Proceedings of S&P*, pages 114–127. IEEE, 1995.
42. A.W. Roscoe and M.H. Goldsmith. What is intransitive noninterference? In *Proceedings of CSF*, pages 228–228. IEEE, 1999.
43. A.W. Roscoe, J.C.P. Woodcock, and L. Wulf. Non-interference through determinism. In *Proceedings of ESORICS*, pages 31–53. Springer, 1994.
44. John Rushby. *Noninterference, transitivity, and channel-control security policies*. SRI International, Computer Science Laboratory, 1992.
45. Peter YA Ryan and Steve A Schneider. Process algebra and non-interference. *Journal of Computer Security*, 9(1):75–103, 2001.
46. A. Sabelfeld and D. Sands. Dimensions and principles of declassification. In *Proceedings of CSFW-18*, pages 255–269. IEEE Computer Society, 2005.
47. Fredrik Seehusen and Ketil Stølen. Information flow security, abstraction and composition. *IET Information Security*, 3(1):9–33, 2009.
48. Geoffrey Smith. On the foundations of quantitative information flow. In *Foundations of Software Science and Computational Structures*, pages 288–302. Springer, 2009.
49. David Sutherland. A model of information. In *Proc. 9th National Computer Security Conference*, pages 175–183, 1986.
50. R. van der Meyden and C. Zhang. A comparison of semantic models for noninterference. *Theoretical Computer Science*, 411(47):4123–4147, 2010.
51. Ron van der Meyden. What, indeed, is intransitive noninterference? In *Proceedings of ESORICS*, pages 235–250. Springer, 2007.
52. J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behaviour*. Princeton University Press: Princeton, NJ, 1944.
53. G. Winskel and M. Nielsen. Handbook of logic in computer science (vol. 4). chapter Models for Concurrency, pages 1–148. Oxford University Press, 1995.
54. J.T. Wittbold and D.M. Johnson. Information flow in nondeterministic systems. In *IEEE Symposium on Security and Privacy*, pages 144–144, 1990.
55. Z. Yin, D. Korzhyk, C. Kiekintveld, V. Conitzer, and M. Tambe. Stackelberg vs. Nash in security games: interchangeability, equivalence, and uniqueness. In *Proceedings of AAMAS*, pages 1139–1146. IFAAMAS, 2010.
56. Aris Zakynthinos and E Stewart Lee. The composability of non-interference. *Journal of Computer Security*, 3(4):269–281, 1995.

57. S. Zdancewic and A.C. Myers. Observational determinism for concurrent program security. In *Proceedings of CSFW-16*, pages 29–43. IEEE Computer Society, 2003.

Appendix

This appendix contains the proofs of some of the propositions and lemmas in the paper.

Proof (Proof of Proposition 2). As the relation R_M^* is constructed by adding related states in several steps, we do the proof by induction on these steps. We show that firstly $R_0 \subseteq \sim_{NI_L}$, and secondly all related pair of states added in each step also is in \sim_{NI_L} .

Induction base: If $(s_1, s_2) \in R_0$, then for some sequence of personalized actions of High players α , either $s_1 \xrightarrow{\alpha} s_2$, or $s_2 \xrightarrow{\alpha} s_1$, hence by property (LR) , and transitivity of the unwinding relation it holds that $(s_1, s_2) \in \sim_{NI_L}$. Therefore $R_0 \subseteq \sim_{NI_L}$.

Induction step: We show that if $R_i \subseteq \sim_{NI_L}$, then $F(R_i) \subseteq \sim_{NI_L}$ holds. $F(R_i)$ is constructed by union of three sets. We show that all these three sets are subsets of \sim_{NI_L} :

i- $R_i \in \sim_{NI_L}$ by the induction step assumption.

ii- If $(s_1, s_2) \in R_i$ then by induction step assumption $(s_1, s_2) \in \sim_{NI_L}$. So for any $t_1, t_2 \in St$, $a \in \mathfrak{A}$ and $l \in L$ such that $do'(s_1, l, a) = t_1$ and $do'(s_2, l, a) = t_2$, by property (SC) of unwinding relation it holds that $(t_1, t_2) \in \sim_{NI_L}$. Therefore:

$$\begin{aligned} & \{(t_1, t_2) \mid \exists (s_1, s_2) \in R, l \in L, a \in \mathfrak{A} \cdot \\ & \quad do'(s_1, l, a) = t_1, do'(s_2, l, a) = t_2\} \\ & \subseteq \sim_{NI_L}. \end{aligned}$$

iii- If $(t_1, s) \in R_i$ and $(s, t_2) \in R_i$, then by induction step assumption it holds that $(t_1, s) \in \sim_{NI_L}$ and $(t_2, s) \in \sim_{NI_L}$. Therefore by transitivity of \sim_{NI_L} it entails that $(t_1, t_2) \in \sim_{NI_L}$, hence:

$$\begin{aligned} & \{(t_1, t_2) \mid \exists s \in St \cdot (t_1, s) \in R \text{ and } (s, t_2) \in R\} \\ & \subseteq \sim_{NI_L}. \end{aligned}$$

By i, ii, and iii we infer that $F(R_i) \subseteq \sim_{NI_L}$, and therefore by induction base and induction step we have that $R_M^* \subseteq \sim_{NI_L}$.

Proof (Proof of Proposition 3). If $NI_M(H, L)$ then by Proposition 1 there is an unwinding relation \sim_{NI_L} on the states of M . By Proposition 2 $R_M^* \subseteq \sim_{NI_L}$ and therefore for any $l \in L$, $s_1, s_2 \in St$ such that $(s_1, s_2) \in R_M^*$ it is the case that $(s_1, s_2) \in \sim_{NI_L}$ and therefore $obs(s_1, l) = obs(s_2, l)$. Hence by Lemma 1 R_M^* is an unwinding relation on the states of M .

Proof (Proof of Lemma 2). Assume $(s_1, s_2) \in R_M^*$ and $l \in L$. For proving that $(obs(s_1, l), obs(s_2, l)) \in \mathcal{U}_M^*$ we verify the conditions in Definition 14. By taking $obs(s_1, l) = obs_1$ and $obs(s_2, l) = obs_2$, conditions (a) and (b) are satisfied trivially. If we take $t_1 := s_2$ and $t_2 := s_2$, then $(s_1, t_1) \in R_M^*$ by the proposition assumption and $(s_2, t_2) \in R_M^*$ by reflexivity of R_M^* . These prove conditions (c) and (d). Condition (e) is also satisfied because $t_1 = t_2$. Therefore $(obs(s_1, l), obs(s_2, l)) \in \mathcal{U}_M^*$.

Proof (Proof of Proposition 4). For the proof, we are going to use Lemma 1 and show that for any two states s_1, s_2 and $l \in L$, if $(s_1, s_2) \in R_{\mathcal{U}_M^*(M)}^*$ then $obs^*(s_1, l) = obs^*(s_2, l)$. First notice that $R_M^* = R_{\mathcal{U}_M^*(M)}^*$, because M and $\mathcal{U}_M^*(M)$ differ only in their observation functions and the definition of R^* relation does not depend on the observation function of the model. So for any $(s_1, s_2) \in R_{\mathcal{U}_M^*(M)}^*$ and $l \in L$ we have that $(s_1, s_2) \in R_M^*$, and by Lemma 2 it follows that $(obs(s_1, l), obs(s_2, l)) \in \mathcal{U}_M^*$, and therefore $obs^*(s_1, l) = obs^*(s_2, l)$. Hence by Lemma 1 it holds that $R_{\mathcal{U}_M^*(M)}^*$ is an unwinding relation for $\mathcal{U}_M^*(M)$ and therefore $NI_{\mathcal{U}_M^*(M)}(H, L)$.

Proof (Proof of Proposition 5). Assume U is a unification of observations for model M such that $NI_{U(M)}(H, L)$ and assume $(obs_1, obs_2) \in \mathcal{U}_M^*$. We show that $(obs_1, obs_2) \in U$ and hence $\mathcal{U}_M^* \subseteq U$. By the definition of \mathcal{U}_M^* , there exists $s_1, s_2, t_1, t_2 \in St$, $l \in L$ such that $obs(s_1, l) = obs_1$, $obs(s_2, l) = obs_2$, $(s_1, t_1) \in R_M^*$, $(s_2, t_2) \in R_M^*$ and $obs(t_1, l) = obs(t_2, l)$. By $NI_{U(M)}(H, L)$ and Proposition 3 we have that $R_{U(M)}^*$ is an unwinding relation for $U(M)$. So as $R_M^* = R_{U(M)}^*$, R_M^* is also an unwinding relation for $U(M)$. Therefore by property (OC) of unwinding relation, from $(s_1, t_1) \in R_M^*$ and $(s_2, t_2) \in R_M^*$ we entail that $obs'(s_1, l) = obs'(t_1, l)$ and $obs'(s_2, l) = obs'(t_2, l)$. Using the definition of $obs'(\cdot)$ we have that $(obs(s_1, l), obs(t_1, l)) \in U$ and $(obs(s_2, l), obs(t_2, l)) \in U$. So, as $obs(t_1, l) = obs(t_2, l)$ and by transitivity property of U , we infer that $(obs(s_1, l), obs(s_2, l)) \in U$, and it follows that $(obs_1, obs_2) \in U$. Therefore $\mathcal{U}_M^* \subseteq U$.

Proof (Proof of Proposition 7). “ \Leftarrow ” Suppose that there exists an unwinding relation \sim_{NI_L} satisfying (OC), (SC) and (LR). We show for all $\alpha \in (\mathfrak{U} \times \mathfrak{A})^*$ and $u_l \in L$, if $exec(\alpha) \neq undef$ then $[exec(\alpha)]_{u_l} = [exec(Purge_H(\alpha))]_{u_l}$. We prove by induction on the size of α .

Induction base: $\alpha = \langle \rangle$. In this case $Purge_H(\alpha) = \langle \rangle$, and therefore $exec(\alpha) = exec(Purge_H(\alpha)) = s_0$. By the reflexivity of \sim_{NI_L} we have that $exec(\alpha) \sim_{NI_L} exec(Purge_H(\alpha))$.

Induction step: Suppose for some $\alpha \in (\mathfrak{U} \times \mathfrak{A})^*$, $exec(\alpha) \neq undef$ implies $exec(\alpha) \sim_{NI_L} exec(Purge_H(\alpha))$. We show that for all $a \in \mathfrak{A}$ and $u \in \mathfrak{U}$ it holds that $exec(\alpha \circ (u, a)) \neq undef$ implies $exec(\alpha \circ (u, a)) \sim_{NI_L} exec(Purge_H(\alpha \circ (u, a)))$ (where \circ denotes the concatenation operator). We consider three cases:
i) If $exec(\alpha \circ (u, a)) = undef$ then it holds that $exec(\alpha \circ (u, a)) \neq undef$ implies $exec(\alpha \circ (u, a)) \sim_{NI_L} exec(Purge_H(\alpha \circ (u, a)))$.

ii) If $exec(\alpha \circ (u, a)) \neq undef$ and $u \in L$ then firstly notice that $exec(Purge(\alpha \circ (u, a))) \neq undef$. Because by induction step assumption and (OC) it holds that $obs(exec(\alpha), u) = obs(exec(Purge_H(\alpha)), u)$ and so because $a \in act(exec(\alpha), u)$, by our model restrictions it holds that $a \in act(exec(Purge_H(\alpha)), u)$. Therefore by $exec(\alpha) \sim_{NI_L} exec(Purge_H(\alpha))$ (induction step assumption), $u \in L$, and (SC) we have $exec(\alpha \circ (u, a)) \sim_{NI_L} exec(Purge_H(\alpha \circ (u, a)))$.

iii) If $exec(\alpha \circ (u, a)) \neq undef$ and $u \in H$ then by (LR) property of \sim_{NI_L} , $exec(\alpha) \sim_{NI_L} exec(\alpha \circ (u, a))$. By this, induction step assumption and $Purge_H(\alpha) = Purge_H(\alpha \circ (u, a))$ we infer that $exec(\alpha \circ (u, a)) \sim_{NI_L} exec(Purge_H(\alpha \circ (u, a)))$.

" \Rightarrow " Suppose that $NI_M(H, L)$, we show there exists an unwinding relation on the states of M . Consider the relation \sim defined as follows: for any $s, t \in St$, $s \sim t$ if for all $\alpha \in (\mathfrak{U} \times \mathfrak{A})^*$ and $u_L \in L$, it holds that if $exec(s, \alpha) \neq undef$ and $exec(t, \alpha) \neq undef$, then $[exec(s, \alpha)]_{u_L} = [exec(t, \alpha)]_{u_L}$. It can easily be seen that \sim is an equivalence relation, we prove that it satisfies (OC), (SC) and (LR) properties.

(OC): If $s \sim t$ and we take $\alpha = \langle \rangle$, by $[exec(s, \alpha)]_{u_L} = [exec(s, \alpha)]_{u_L}$ it holds that $[s]_{u_L} = [t]_{u_L}$ and therefore \sim satisfies (OC).

(SC): Suppose that for some $s, t \in St$, $u \in L$ and $a \in \mathfrak{A}$ such that $s \sim t$, it holds that $do(s, u, a) \neq undef$ and $do(s, u, a) \not\sim do(t, u, a)$. Then there exists $\alpha \in (\mathfrak{U} \times \mathfrak{A})^*$, $u_L \in L$ such that $exec(do(s, u, a), \alpha) \neq undef$, $exec(do(t, u, a), \alpha) \neq undef$, and $[exec(do(s, u, a), \alpha)]_{u_L} \neq [exec(do(t, u, a), \alpha)]_{u_L}$. Therefore $[exec(s, ((u, a) \circ \alpha))]_{u_L} \neq [exec(t, ((u, a) \circ \alpha))]_{u_L}$, which contradicts $s \sim t$.

(LR): Suppose that for some $s \in St$, $u \in H$ and $a \in \mathfrak{A}$, it holds that $do(s, u, a) \neq undef$ and $s \not\sim do(s, u, a)$. Then there exists $\alpha \in (\mathfrak{U} \times \mathfrak{A})^*$, $u_L \in L$ such that $exec(do(s, u, a), \alpha) \neq undef$, $exec(s, \alpha) \neq undef$, and $[exec(s, \alpha)]_{u_L} \neq [exec(do(s, u, a), \alpha)]_{u_L}$. Because s is reachable, we have that $s = exec(\beta)$ for some $\beta \in (\mathfrak{U} \times \mathfrak{A})^*$. Therefore $[exec(\beta \circ \alpha)]_{u_L} \neq [exec(\beta \circ (u, a) \circ \alpha)]_{u_L}$. But this is a contradiction because by $NI_M(H, L)$ it holds that $[exec(\beta \circ (u, a) \circ \alpha)]_{u_L} = [exec(Purge_H(\beta \circ (u, a) \circ \alpha))]_{u_L}$ and $[exec(\beta \circ \alpha)]_{u_L} = [exec(Purge_H(\beta \circ \alpha))]_{u_L}$ and we have that $Purge_H(\beta \circ (u, a) \circ \alpha) = Purge_H(\beta \circ \alpha)$.